

Remarks

Claims 7-8 (Currently amended)

Claims 7-8 of Application have been amended to depend solely on Claim 1. This should remove objection under 37 CFR 1.75(c).

In regards to any possible conflict or anticipation by Patent 6,516,350 (Lumelsky et al.) hereafter referred to as "Lumelsky"), while we recognize that Claims 7-8 of subject Application 10/058,242 (hereinafter referred to as "Application") have yet to be evaluated, we wish to state that Lumelsky does not appear to address "Peer-to-Peer" operation for storage distribution and usage in any way, and we thus anticipate no conflict. Likewise, Lumelsky contains no reference to usage within a network switch device, and thus we expect no conflict with Claim 8.

As to the rejection of Claims 1-6 under 35 USC 102, on the grounds of anticipation by Patent 6,516,350 (Lumelsky et al.):

Claim 1 (Currently amended)

The intention of Lumelsky is to teach a method of replicating content for multimedia delivery, and, inter alia, it describes a means for controlling access and flow to this multimedia content. While other uses might be construed by inference, Lumelsky talks to multimedia usage in many places, and offers multimedia as the preferred embodiment. The method taught defines a means to find a set of copies of the streaming content, available on multiple servers, and then determine which copy is the best one to use. This results in more customers receiving streamed data.

The subject Application addresses different technical problems, and a different part of the computer structure, to Lumelsky. A network of (multimedia) distributed computer resources, as described by Lumelsky, requires a means of determining which server has the requested content.

Not discussed by Lumelsky, is that, within that server, and the other servers in the network, there resides a multiplicity of file-systems, which describe where data is stored on the storage media. Where any data that is stored on any given media can be accessed by multiple servers, it is necessary to coordinate the file-system copies kept in each server to prevent changes by one server corrupting operations in another server, a practice known as maintaining file-system integrity.

In contrast to Lumelsky, the focus of the Application is the maintenance of file-system integrity for any computer files accessible across a set of Network-Attached Storage (NAS) computers, while obviating the time-consuming process required to maintain exact and instantaneous duplication of the file-system directory structure on each and every server (file-system integrity) that is created with, so-called, global file systems. This is done by apportioning the total file system for the stored data in such a way that each NAS computer contains part of the file system, selected in such a way as to give roughly even workload to each NAS computer. This results in a lower usage of system resources, and a much lower time to change file structure, when compared

with current means that update a file map in every single NAS computer simultaneously, or that use a single 'lock-box' copy of the file system, as shown in Figures 1,2 of Application.

Lumelsky (and specifically Col9, lines 15-30) does not address the file-system issue. It focuses on means to find and address a single content server from a set of alternatives, with the means residing in a client computer or in a directory server system in the network. The Application teaches how the file-systems in the NAS or content servers operate, and provides a mechanism that potentially complements Lumelsky in actual usage, but which addresses a different technical area with a different mechanism for operation.

Claim 1 Para a) and b) are amended herewith for clarity on this issue. As amended, they are not anticipated by Lumelsky (and specifically Col9, lines 15-30), following the distinction drawn above that Lumelsky does not focus on the file-system.

Claim 1 Para c) as amended also clearly addresses file-system issues, and Lumelsky Col 10, lines 18-44 can now be seen as not anticipatory, since it does not describe file-systems. Further, Lumelsky, col10, lines 18-44 does not describe a "specific designation as the initial contact point" process.

Claim 1 Para d) stands within the context of file-systems. An initial contact point, followed by redirection, is a common means for accessing networked resources, predating Lumelsky. With the location of the "Management Plane" in Fig 4 of Lumelsky being undefined, it would be reasonable to compare this method as taught by Lumelsky with the method used since the early 1990's by Microsoft in the Windows Operating System to access network files. In the Application, this initial contact point is one of the set of computer servers, or a separate dedicated server, and so the Microsoft comparison is not appropriate.

In any event, the comparison must be made in the light of the different focus of Lumelsky, compared with the Application.

Claim1 Para e) reflects a means to add more computer elements to the network. It describes the re-partition of the file-system ownership over the extended network, and likewise, if a computer element fails or is removed. Lumelsky, Col 11, lines 31-59 addresses the replication of the data across servers, rather than a repartition of the file-system control mechanism. These are fundamentally different processes. To help clarify this, Claim 1Para e) has been amended.

With the above comments, and amendments as presented, we respectfully believe that Patent 6,516,350, Lumelsky et al. does not anticipate Claim 1 of the Application.

Claim 2 (Currently amended)

In Lumelsky, Col 11, lines 31-59, means are described for replication of data objects. However, said section of Lumelsky does not describe the purpose of this as 'backup'. Rather, in line 46, and again in line 58, the purpose is described as "shaping capacity" or "incrementing or decrementing capacity".

On this ground alone, we submit that Lumelsky does not anticipate Claim 2 of subject Application.

However, Lumelsky, as described above for Claim 1, attacks a different technical problem and are to Application. We have amended the phrasing of Claim 2 to better reflect this, and to highlight that replication for backup as taught in Application is in fact replication of control information and processes, and not of the data objects themselves as taught by Lumelsky.

On this ground, too, we submit that Lumelsky does not conflict with or anticipate Claim 2 of Application.

Claim 3 (Original)

Likewise to claim 2, Lumelsky describes replication of objects only. In contrast, Claim 3 of Application teaches a means to provide multiple computer elements access to a given data object, by mapping replicas of the control information (not data objects) in such a way that the essence of file-system integrity is maintained.

We submit that Lumelsky does not anticipate or conflict with Claim 3 of application, based on having a fundamentally different mechanism, and on working in different parts of the network system.

Claim 4 (Original)

As discussed above regarding Claim 2, replication of data objects in the Lumelsky patent is described as being for “capacity” purposes. This is further borne out in the Claims Section of Lumelsky, where there is no reference to “backup”, “failure protection” or such similar phrases regarding prevention of the loss of data availability as described in Application. Irrespective of whether the methods used to achieve replication are different in embodiment, we submit that usage of replication in Application is different from that of Lumelsky, and that therefore Lumelsky does not anticipate subject Application as regards Claim 4.

Claim 5 (Currently amended)

The phrase “two-tier” has been removed in the attached Amendment to Claims for subject Application. This is solely to avoid confusion with the Lumelsky description of two layers. Lumelsky nowhere describes the management of empty or ‘free’ space on the storage media. This is a fundamentally different process to data object replication, which, as mentioned above, is addressed by Lumelsky.

The means described in Application, or an alternative, are required whenever some part of the total storage media can be accessed by more than one computer element. Claim 5 reflects a mechanism, within the means described in the Application, for efficiently distributing the unused, free space, so that each computer element can provide for the expansion of files as more data is added, while retaining the bulk of the unused, free space for allocation, by request, to any computer element that uses up its local allocation.

This type of mechanism, or this usage, is lacking from Lumelsky, and we submit that Claim 5 of subject Application is neither anticipated by, or in conflict with, Lumelsky.

Claim 6 (Original)

While Lumelsky describes negotiation for optional services involving encryption of data objects, and user authentication and authorization (Col 10, lines 45-65), neither here nor in the Claims (Col 18, lines 4-12), does it describe provision of these services. Application not only teaches this as a means, but also extends the concept of encryption beyond authentication to cover the file-system and all the communications between client systems and the storage systems.

We submit that these are fundamentally different, that Lumelsky does not in act describe an encrypted system, but only the mechanism to connect to one if it existed, and that the lsky description does not encompass the usage described in subject Application's Claim 6.

We therefore submit that Lumelsky does not anticipate Claim 6 of the Application.

Conclusion

Claims 1-6 of the Application are amended to more clearly reflect the focus of Application on the file-system. With the different focuses as described above, and with this amendment, we therefore respectfully submit that Lumelsky does not anticipate Claims 1-6 of Application

FOR INFORMATION ONLY

This page and the following 3 pages are “clear-copy” versions of the Claims in Application 10/058,242, as amended 23rd July 2005 and further corrected Sep 16, 2005 to show Claim status, and are intended for informational use only.

- 1) (Currently amended) A means of building a Scalable Network-Attached Storage system where control of the data storage elements of said system is distributed over the computer elements of said system, so allowing said computer elements to access and control said data elements in a shared fashion whereby:
 - a) Control information describing the physical location on storage media of the blocks of data stored on said storage media and describing the file-system structure determining the inter-relationship of said blocks of data that form a set of data storage elements can reside in different computer elements, together with the control processes necessary to read, write, add, delete or otherwise change such control information, so allowing large numbers of computer elements to be used in the Scalable Network-Attached Storage system, and so allowing said system to be easily and economically expanded in size and performance and reconfigured to needs;
 - b) Said control information for any such data storage element can be replicated in several computer elements in such a way that these several computer elements can access said data elements;
 - c) Allocation to the set of computer elements of said control of said data elements and their physical sub-elements is initially established by a software functionality according to a set of user and computer generated policy rules; and where said software functionality adjusts said distribution of control across said computer elements on a periodic basis, depending on metrics measured periodically throughout the Scalable Network-Attached Storage system
 - d) Computer elements are specifically designated as the initial contact point for a client computer to the Scalable Network-Attached storage system, which designated computer elements have a software facility to determine that computer element having control of the data storage element that said client computer wishes to access and by means of said

software facility re-direct said client computer to communicate directly with that computer element;

- e) Where the software facility of paragraph c) in Claim 1) above detects the addition of new computer elements to the Scalable Network-Attached Storage system via the periodic metrics transmitted to said software facility by said new computer element; and thereby said software facility re-maps the allocation to said computer elements of the control information and control processes of said data storage elements to make use of said new computer element; and where the loss of a computer element through failure or removal is detected by a loss of periodic metric data, so causing said software facility to re-map the allocation to the remaining computer elements of the control information and control processes of those data storage elements previously controlled by said lost computer element;

- 2) (Currently amended) An extension of the means of Claim 1) where temporary loss of access to said data storage elements is reduced by having a prepared second copy of the control information and control processes of the Scalable Network-Attached Storage system, whereby another computer element, which is configured to rapidly take control of said data elements, is designated as backup for each computer element in such a way that a hardware or software failure will not affect both said computer element and its designated backup simultaneously; so that a failure detected by the software facility in Para e) of Claim 1) will cause that software facility to move control to said backup computer element;

And where said backup computer element may be a computer element or one of a set of computer elements specifically and solely functioning as backup computer elements;

Or where said backup computer element may be a computer element that is actively controlling access to other data storage elements.

- 3) (Original)An extension to the means of Claim 1) where the performance of the Scalable Network-Attached Storage system is increased by mapping multiple computer elements to be able to control any given data storage element;

Where one of said controlling computer elements is designated as the sole computer element allowed to change said data storage element, with the other computer elements being able to read said data storage element; or where the type of data storage element or the client

computers' method of accessing same permits multiple computer elements to change said data storage element.

- 4) (Original) An extension to Claim 3 whereby the ability of the system to recover from a failure is enhanced by having a data replication software facility which allows a computer element to replicate data according to policy rules managed by the system, with said replication being to both local computer elements who store such replicas on the data storage elements under their control, or to remote computer elements, permitting copies of data to be at a safe distance to protect against natural or man-made disasters;

And where said data replication software facility may also be used to provide copies of data at the remote site or sites for said remote site or sites to be able to access data more rapidly than if it were at the originating site;

And where the said policy rules for replication may include schedule, frequency and priority of replication, number of backup copies, type of backup data storage elements and other policy rules.

- 5) (Currently amended) An extension of Claim 1 above whereby a system is used to manage unused free space in the Scalable Network-Attached Storage system and its derivatives, such system being implemented as a software facility that provides both a means to allocate part of the available unused free space to each computer element, keeping the remainder under its own control, and which software facility uses policy rules to monitor, control and change this allocation periodically based on metric information reported to said software facility by the computer elements.

- 6) (Original) An extension to the means in the above claims whereby a set of the data storage elements and computer elements in a Scalable Network-Attached Storage system are designated as a Secure Scalable Network-Attached Storage system, with the data storage elements being encrypted by the client computer, and with the file structure of said data elements being encrypted, and with communications between the client systems and the Scalable-Network-Attached storage system being encrypted.

- 7) (Currently amended) An extension to the means in claim 1 whereby those means are employed to take advantage of data storage elements in the client computers by using said data storage elements in part or in whole as data storage elements in a distributed form of Scalable Network-Attached Storage (here named Peer-Based Storage Network), where

- a) The data storage elements in any given client computer can be shared with other client computers under the control of the Scalable Network-Attached Storage software forming this invention, as extended to provide the Peer-Based Storage network capability, and with the features of replication and security as claimed above and described herein;
 - b) Where software facilities are provided to allow designation of the amount of said data storage elements a client computer may wish to share; with software facilities to increase or decrease said amount as desired without causing loss of data to the client computers sharing the data storage elements being changed;
 - c) Where replication and backup policies may be developed to protect automatically against loss of a client computer and/or its stored data on its data storage elements
- 8) (Currently amended) An extension to the means in Claim 1 of the invention whereby the computer elements of the Scalable Network-Attached Storage system, Secure Scalable Network-Attached Storage system or Peer-Based Storage Network in part or in whole are contained within a storage network switching element or a local area network switching element or a communications switching element.